arXiv:1612.04746v1 [cs.GT] 14 Dec 2016

A Simple and Approximately Optimal Mechanism for a Buyer with Complements

Alon Eden* Michal Feldman† Ophir Friedler‡ Inbal Talgam-Cohen§
S. Matthew Weinberg¶

December 15, 2016

Abstract

We consider a revenue-maximizing seller with m heterogeneous items and a single buyer whose valuation v for the items may exhibit both substitutes (i.e., for some S, T, $v(S \cup T) < v(S) + v(T)$) and complements (i.e., for some S, T, $v(S \cup T) > v(S) + v(T)$). We show that the better of selling the items separately and bundling them together guarantees a $\Theta(d)$ fraction of the optimal revenue, where d is a measure on the degree of complementarity. Note that this is the first approximately optimal mechanism for a buyer whose valuation exhibits any kind of complementarity, and extends the work of Rubinstein and Weinberg [2015], which proved that the same simple mechanisms achieve a constant factor approximation when buyer valuations are subadditive, the most general class of complement-free valuations.

Our proof is enabled by the recent duality framework developed in Cai et al. [2016], which we use to obtain a bound on the optimal revenue in this setting. Our main technical contributions are specialized to handle the intricacies of settings with complements, and include an algorithm for partitioning edges in a hypergraph. Even nailing down the right model and notion of "degree of complementarity" to obtain meaningful results is of interest, as the natural extensions of previous definitions provably fail.

*Computer Science, Tel-Aviv University. alonarden@gmail.com. Work done in part while the author was visiting the Simons Institute for the Theory of Computing.

†Computer Science, Tel-Aviv University, and Microsoft Research. michal.feldman@cs.tau.ac.il.

‡Computer Science, Tel-Aviv University. ophirfriedler@gmail.com.

§Computer Science, Hebrew University of Jerusalem. inbaltalgam@gmail.com.

¶Computer Science, Princeton University. smweinberg@princeton.edu. Work done in part while the author was a research fellow at the Simons Institute for the Theory of Computing.

1 Introduction

Consider a revenue-maximizing seller with m items to sell to a single buyer. When there is just a single item, and the buyer's value is drawn from some distribution with CDF F, seminal works of Myerson [36], and Riley and Zeckhauser [39] prove that the optimal mechanism is to simply set whatever price maximizes $p \cdot (1 - F(p))$. It soon became well-understood that beyond the single-item setting, the optimal mechanism suffers many undesireable properties that make it unusable in practice, including randomization, non-monotonicity, and others [40, 28, 29, 5, 12, 13, 43, 38]. Following seminal work of Chawla, Hartline, and Kleinberg [8], there is now a sizeable body of research proving that the simple mechanisms we see in practice are in fact approximately optimal in quite general settings, helping to explain their widespread use [9, 10, 32, 27, 2, 3, 41, 34, 44, 7, 11].

Still, prior work has largely been limited to additive[1] or unit-demand[2] buyers. Only recently have researchers begun tackling more complex valuation functions, and even these works have remained restricted to subclasses of *subadditive* valuations, also called *complement-free* [41, 11, 7].[3] While subadditive valuations are quite general, they can only capture interaction between items as *substitutes*. For example, if the items are pieces of furniture, a buyer's marginal valuation for a chair might decrease as her home gets more and more filled due to lack of space. To date, no results in this line of work model iteraction between items as *complements*. For example, a buyer's value for a kitchen table might actually increase if she already has a chair to sit. The goal of this paper is to study simple and approximately optimal mechanisms in domains (like the example above) where buyer valuations exhibit both substitutes and complements.

Buyers with Complements. Even for the traditionally simpler domain of *welfare* maximization, the state-of-the-art only recently has begun designing mechanisms for buyers with complements [1, 20, 24, 22]. The main difficulty is that horrible lower bounds are known for general valuations [37], so in order to get interesting positive results, some assumptions are necessary on the degree to which buyer valuations exhibit substitutes or complements. Interestingly, good positive results are possible in the complete absence of complements and no restriction on the degree of substitutability [17, 16, 19, 23, 15], but not vice versa: many strong lower bounds still exist in the absence of substitutes but with arbitrary complementarity [33, 1, 35, 22].

So the goal of these recent works is to parameterize the "degree of complementarity" that a valuation function admits, and prove an approximation guarantee of $f(d)$ whenever buyer valuations have "complementarity of degree at most d" [1, 21, 20, 24, 22]. For example, if you were selling a table, chair, bicycle, banana, and socks, you would reasonably expect buyers to view the table and chair as complements, but likely not the bicycle and banana. Similarly, you wouldn't expect any set of three items to be viewed as complements (outside of what's already captured by the table and chair as a pair). So it seems overly pessimistic not to try and exploit this. Ideally, a good formal definition for "complements of degree d" should make sense in its own right (i.e. without appealing to results) and capture a smooth transition as d grows (i.e. we don't have $f(0) = 1$ and $f(d) = m$ for all $d > 0$). Interestingly, the right formal definition of "complements of degree d" seems to differ between environments. Some examples of previous successful definitions include the "supermodular degree" and "positive-hypergraphs degree" [21, 20].

Nailing down the right model of complementarity degree is even trickier for the revenue objective, as we must also fold some notion of independence into the value distribution in order to avoid

[1] A buyer valuation is additive if $v(S) = \sum_{i \in S} v(\{i\})$.
[2] A buyer valuation is unit-demand if $v(S) = \max_{i \in S}\{v(\{i\})\}$.
[3] A valuation is subadditive if $v(S \cup T) \leq v(S) + v(T)$ for all S, T.

extremely strong lower bounds that hold against even additive valuations over two items [5, 28].[4] We postpone a formal definition of our model and corresponding notion of complementarity degree to Section 2, and give an illustrative example here. Imagine you are selling furniture and related goods. Some items naturally exhibit complementarities: with a table, chair, and silverware, a buyer can eat meals at home. With a table, four chairs, and a game of Settlers, they could host a board games night. With two sofas and a TV, they could instead host a movie night. So think of the buyer as having a non-negative valuation for being able to eat meals in their home, host events, etc., and their preliminary value for a set S of items is additive over the activities that S allows them to partake in.[5] But there's a catch: no buyer has room in their apartment to comfortably fit a table, TV, two sofas, and four chairs. So the items are also substitutes - even if the buyer were to wind up with the entire warehouse of furniture, they won't get use out of anything besides what fits in their apartment. So we let $\mathcal{C}$ represent a set system determining which items can fit in the apartment, and let $w(T)$ denote the buyer's value for whatever special activity the items in exactly T allow her to partake in (that she couldn't partake in with any proper subset). The buyer's value for a set of items S is the maximum over all S' that fit in her apartment of the sum of her values for all the activities she can partake in using S'. Independence enters the model by assuming buyers have independent values for different activities, and the degree of complementarity is captured via the maximum number of activities that require any given item.

Main Result. Our main result (Theorem 3.2) is that the better of selling separately (post a price on each item, let the buyer purchase whatever subset she likes) or bundling together (post a single price on the grand bundle, let the buyer purchase or not) achieves a tight $\Theta(d)$ approximation whenever buyer valuations exhibit complementarity at most d.

We also complete the picture by showing that our notion of complementarity is in some sense the right one: if instead we measure complementarity via the "supermodular degree," then there exist populations in our model with supermodular degree d for which the better of selling separately and bundling together achieves only a $\Omega(2^d/d)$-approximation. Similarly, if we instead measure complementarity via the "positive-hypergraph degree," then there exist populations in our model with positive-hypergraph degree d for which the better of selling separately and bundling together achieves only a $\Omega(\sum_{\ell \leq d} \binom{m}{\ell}/m)$-approximation. Both notions of degree are defined formally in Section 6 where the lower bounds are proved. The point is not that $\Theta(d)$ is a "better" bound than $\Omega(2^d/d)$, as this is in some sense not a fair comparison, but rather that "supermodular degree" and "positive-hypergraph degree" are incapable of capturing the smooth transition from low degrees to high degrees of complementarities as they can only take on m different values but provide guarantees that range from 1 to $\Omega(2^m)$. In comparison, our notion of degree of complementarity takes on 2^{m-1} different values, and provides guarantees that range from 1 to $\Omega(2^m)$, allowing for an exponentially finer-grained tradeoff.

Our Techniques. Our starting point is a duality-based upper bound on the optimal achievable revenue coming from recent work of [6]. Their upper bound decomposes into three parts, which they call SINGLE, CORE, and TAIL. So the goal is to show that selling separately well-approximates SINGLE, and that bundling together well-approximates CORE and TAIL. Fortunately, the analysis of [6] is fairly robust, and we are able to prove that bundling together achieves a constant factor of both CORE and TAIL via a similar approach. Our main technical contribution appears

[4] Specifically, there exists a distribution D over $\mathbb{R}^2$ such that when a single additive buyer's valuation is drawn from D, the optimal revenue for the seller is infinite, but the revenue of the best deterministic mechanism is 1.

[5] Such valuations functions are called "positive-hypergraph" (PH) valuations.

in Section 4, where we prove that selling separately gets an $O(d)$-approximation to SINGLE. Incidentally, bounding SINGLE happened to be the easiest part of the analysis in [6] for additive valuations.

Without getting into details about what exactly this SINGLE term is, we can still highlight the key challenge. Essentially, we would like to post a different price on each activity. In fact, we can show that the optimal "activity-pricing scheme" even obtains a constant-factor approximation to SINGLE. The catch is that we sell items, not activities. We may wish to set drastically different prices on many different activities requiring the same item, and it's unclear that we can achieve the desired activity prices by cleverly setting prices on the items separately (in fact, it could be impossible). So our main technical contribution is an algorithm to find a subset of activities S for which it is possible to achieve any desired activity-pricing on S by only posting prices on items, and the optimal revenue from activities in S is a d-approximation to the optimal *activity*-pricing scheme. It turns out that the right sets of activities to search for are ones where each activity requires an item *not* required by any of the others, that the number of collections with this property necessary to partition all activities tightly characterizes the approximation guarantee of selling separately, and that d collections suffice whenever each item is required by at most d activities.

1.1 Related Work

Multi-Dimensional Auction Design. A rapidly growing body of recent literature has shown that simple mechanisms are approximately optimal in quite general settings [8, 9, 10, 32, 27, 34, 2, 44, 41, 3, 11]. Of these, the result most related to ours is [41], which proves that the better of selling separately and bundling together achieves a constant-factor approximation for a single buyer whose valuation is drawn from a population that is "subadditive with independent items". Their model is similar to our model with $d = 1$ (but neither subsumes the other), so our results can best be interpreted as an extension of theirs to buyers whose valuations also exhibit complementarity.

In terms of techniques, our work makes use of a recent duality framework developed in [6]. The same duality framework has been used in concurrent work by the present authors to prove multi-dimensional "Bulow-Klemperer" results [18], and independent work by others to design simple, approximately optimal auctions for multiple subadditive bidders [7]. Still, the duality theory is only used to provide an upper bound on the revenue in all these cases, and the remaining technical contributions are disjoint. In particular, for the present paper, Section 3 has a high technical overlap with these works, and Section 5 bears some similarity. But our main technical contribution lies in Section 4, which is unique to the problem at hand.

Agents with Complements. In recent years there has also been a rapid growth in the design of algorithms and mechanisms in the presence of complements [1, 21, 25, 26, 20, 24, 22]. These works consider many different aspects: for example, assuming strategic behavior of agents (or not), assuming the existence of strict substitutes (or not), or focusing on simple mechanisms and quantifying the efficiency of equilibria. In all these works, some notion of degree of complementarity was cast on a class of valuation functions, and the approximation ratio guaranteed grew as a function of complementarity degree. It is noteworthy that quite often different settings motivate different degrees of complementarity to best capture the degradation in possible guarantees. For instance, [1] uses the positive hypergraph (PH) degree, [21] uses the supermodular degree, [20, 24] use the maximum over PH degree, and [22] uses the positive supermodular degree.

In comparison to this literature, ours is the first to consider *revenue* maximization for buyers with complements.

1.2 Discussion and Future Work

We present the first simple and approximately optimal mechanism for a buyer whose valuation exhibits both substitutes and complements. We show that for a natural notion of "degree of complementarity," the better of selling separately and selling together achieves a tight $\Theta(d)$-approximation to the optimal revenue. We provide rigorous evidence that this is "the right" notion to consider via large lower bounds for classes of valuations that previous definitions would deem simple.

Our main technical contribution is an algorithm to partition a collection of sets into subcollections such that each set (in the subcollection) contains an item not contained in the others (in that same subcollection). Due to the robustness of previously-developed tools like the "core-tail" decomposition [34, 2, 41, 44, 11], and duality-based benchmarks [6], we are able to focus our technical contributions to the specific problem at hand.

The obvious direction for future work would be to see whether simple mechanisms remain approximately optimal for *multiple* buyers with complementarity degree d. Doing so would likely require at least one substantial innovation beyond the ideas in this paper, as even the $d = 1$ case remains open (even considering the recent breakthrough result of [7]). Our work also contributes to the growing body of evidence that our community now has the tools to "catch up" the state-of-the-art for multi-dimensional mechanism design to the wealth of knowledge that currently exists for single-dimensional settings. Considering buyers with complements is one important path in this direction, but there are numerous others as well.

2 Preliminaries

Buyer Valuations. We consider a setting in which a seller wishes to sell a set M of m items to a single buyer. The buyer has a valuation function v that assigns a non-negative real number $v(S)$ to every bundle of items $S \subseteq M$. The valuation is normalized ($v(\emptyset) = 0$) and monotone ($v(S) \leq v(T)$ whenever $S \subseteq T$). We also abuse notation and let $v(X) = \mathbb{E}_{S \leftarrow X}[v(S)]$ when X is a random set.

Complementarities. An increasingly popular model to represent complementarities is via a *positive hypergraph representation*. That is, $w : 2^M \to \mathbb{R}^+$ is a non-negative function, and $w(T)$ denotes the bonus valuation that the consumer enjoys from exactly the set of items T (in addition to the value the consumer already enjoys for proper subsets of T), i.e., $v(S) = \sum_{T \subseteq S} w(T)$. In the language of Section 1, $w(T)$ denotes the bidder's value for the activity requiring exactly items in T. We will sometimes refer to T as a *hyperedge*, thinking of $w(\cdot)$ as weight function on the hypergraph with nodes M. As an example, if v is additive, then defining $w(\{i\}) = v(\{i\})$ and $w(T) = 0$ whenever $|T| > 1$ yields $v(S) = \sum_{T \subseteq S} w(T)$. We say that v (or w) exhibits *complementarities of degree d* if for all i, $|\{S \ni i : w(S) > 0\}| \leq d$.

Substitutes. An equally popular model to represent substitutes is via combinatorial constraints. Let $\mathcal{C} \subseteq 2^M$ denote a downwards-closed set system on M. $S \notin \mathcal{C}$ denotes that at least some items in S are substitutes, and the buyer does not derive value from all of S. Many valuations that exhibit *only* substitutabilities are "additive subject to constraints $\mathcal{C}$": $v(S) = \max_{T \subseteq S, T \in \mathcal{C}} \{\sum_{i \in T} v(\{i\})\}$. For example, unit-demand valuations can be represented with $\mathcal{C} = \{T : |T| \leq 1\}$.

Complements and Substitutes. We choose to model substitutes and complementarities together by combining the above two models. That is, there is a positive hypergraph representation w that represents complementarities, and combinatorial constraints $\mathcal{C}$ that represent substitutabilities, and $v(S) = \max_{T \subseteq S, T \in \mathcal{C}} \{\sum_{U \subseteq T} w(U)\}$. Recall the furniture example: w represents that two

sofas and a TV allow you to host a movie night, whereas any proper subset doesn't. $\mathcal{C}$ represents that you can only fit so much furniture in your apartment. We assume w.l.o.g. that $w(T) = 0$ for all $T \notin \mathcal{C}$, as the bidder will never be able to partake in activity T no matter what (because the required items don't fit in the apartment).

Value Distributions. We model our buyer valuation $v(\cdot)$ as being drawn from the population D in the following way. There are some constraints $\mathcal{C}$, that are fixed (not randomly drawn). Each $w(T)$ is then drawn independently from some distribution D'_T for all T, and $v(S) = \max_{T \subseteq S, T \in \mathcal{C}} \{\sum_{U \subseteq T} w(U)\}$. We say that D has complementarity d if all v in the support of D have complementarity d. Note that this implies D has complementarity d if and only if for all i, $|\{T \ni i : \Pr[w(T) = 0] < 1\}| \leq d$. We use V to denote the support of D, $f(v)$ to denote $\Pr_{\hat{v} \leftarrow D}[\hat{v} = v]$, and $f_T(y) = \Pr_{x \leftarrow D'_T}[y = x]$.

Truthful Mechanisms and Revenue Maximization. Formally, a mechanism $\mathcal{M}$ has two mappings $X : V \to \Delta(2^M)$, and $p : V \to \mathbb{R}$. X takes as input a valuation v and awards a (potentially random) subset of items. p takes as input a valuation v and charges a price. $\mathcal{M}$ is then *truthful* if for all $v, v' \in V$, $v(X(v)) - p(v) \geq v(X(v')) - p(v')$.[6] Alternatively, one can view a mechanism as a *menu* that lists options of the form (X, p), where $X \in \Delta(2^M)$ and $p \in \mathbb{R}$. A buyer with value $v(\cdot)$ then selects the menu option $\arg\max\{v(X) - p\}$. It is easy to see the equivalence between the two representations: simply setting $(X(v), p(v)) = \arg\max\{v(X) - p\}$ takes one from the menu view to a truthful mechanism. We denote by $\text{REV}(D)$ the optimal revenue attainable by any truthful mechanism when buyer valuations are drawn from the population D.

Simple Mechanisms. The two simple mechanisms we study are selling separately (SREV) and bundling together (BREV). We denote by $\text{BREV}(D)$ the optimal expected revenue attainable by selling all items together, and will drop the parameter D when it is clear from context. It is well-known that $\text{BREV}(D) = \max p \cdot \Pr[v(M) \geq p]$ [36]. SREV is a touch trickier, as it is NP-hard for buyers in our model to even decide what set of items they wish to purchase at a given set of prices, so it's not even clear how we should evaluate the "revenue" of a price vector. We cope with this using a similar approach to [41]: we define SREV^* to be the optimal revenue attainable by any item pricing *only counting an item as sold if every set the buyer is willing to purchase contains that item.* More formally, for a given item pricing $\vec{p}$, and valuation v, let $P_i(\vec{p}, v) = 1$ if $\exists S \ni i, v(S) - \sum_{j \in S} p_j > 0$ and $\forall S \not\ni i, v(S) - \sum_{j \in S} p_j \leq 0$, and $P_i(\vec{p}, v) = 0$ otherwise. Then $\text{SREV}^*(D) = \max_{\vec{p}} \mathbb{E}_{v \leftarrow D}[\sum_i P_i(\vec{p}, v) \cdot p_i]$.

Discrete vs. Continuous Distributions. Like [6], we only explicitly consider distributions with finite support. Like their results, all of our results immediately extend to continuous distributions as well via a discretization argument of [14, 41, 31, 30, 4]. We refer the reader to [6] for the formal statement and proof. Theorem 2.1 assumes that for every single-dimensional random variable X and number $q \in [0, 1]$, there exists a threshold p so that $X \geq p$ with probability *exactly* q, which might a priori seem problematic for discrete distributions. Fortunately, standard "smoothing" techniques allow this assumption to be valid for discrete distributions. A formal discussion of this appears in Remark 2.4 of [41].

The Copies Environment. In our bounds, we'll make use of a related "copies environment" [8, 9, 10, 32]. For any product distribution $D' = \times_{i=1}^{k} D'_i$, we define the corresponding copies setting

[6] Note that for a single buyer, there is no need to distinguish between Bayesian Incentive Compatible and Dominant Strategy Incentive Compatible - they're the same.

as follows: there is a single item for sale, and k buyers. Buyer i's value for the item is drawn from the distribution D'_i. For instance, in our model, the hypergraph representation of the valuation is drawn from $D' = \times_S D'_S$, so we would have a bidder for every subset, with bidder S's value drawn from the distribution D'_S.

We can then define the benchmark $\text{OPT}^{\text{copies}}(D')$ to be the expected revenue obtained by the optimal mechanism (Myerson's [36]) on input D'. Note that this is equal to $\mathbb{E}_{w \leftarrow D'}\left[\max_T\{\bar{\varphi}_T(w(T)), 0\}\right]$, where $\bar{\varphi}_T(\cdot)$ denotes Myerson's *ironed virtual value* for the distribution D'_T. We will make use of the following theorem from [9]:

Theorem 2.1 ([9])**.** *For any $q \leq 1$, there exist (possibly random) prices $\{p_T\}_T$ such that:*

1. *Revenue is high: $OPT^{copies}(D') \leq \frac{1}{q}\sum_{T\subseteq M}\mathbb{E}_{p_T}\left[p_T \cdot \Pr_{x\leftarrow D'_T}[x \geq p_T]\right]$.*
2. *Probability of sale is low: $\sum_{T\subseteq M}\mathbb{E}_{p_T}\left[\Pr_{x\leftarrow D'_T}[x \geq p_T]\right] \leq q$*
3. *Moreover, each p_T takes on at most two values. If D'_T is regular, then p_T is a point-mass.*[7]

3 Our Duality Benchmark and Main Theorem Statement

We extend the duality framework of [6] to our setting in a natural manner. Full technical details are deferred to Appendix A. The only technical detail needed for stating our revenue benchmark is the following: we partition the valuation space V into $2^m - 1$ different regions, depending on which hyperedge is the most valuable to a buyer with valuation v. Specifically, we say that v is in region R_A if $A = \arg\max_{T\subseteq M}\{w(T)\}$, with ties broken lexicographically.

Corollary 3.1. *For valuation distribution D established by drawing a hypergraph representation $w \leftarrow \prod_S D'_S$ and returning $v(S) = \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T} w(U)\}$.*

$$REV(D) \leq \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[v \notin R_T\right]\}\right] \qquad (NON\text{-}FAVORITE)$$

$$+ \mathop{\mathbb{E}}_{v\leftarrow D}\left[\sum_{S\subseteq M}\max\{0, \bar{\varphi}_S(w(S))\}\cdot\mathbb{1}\left[v \in R_S\right]\right] \qquad (SINGLE)$$

In Section 4, we show that $\max\{\text{SREV}^*, \text{BREV}\}$ gets a $4(d+1)$-approximation to SINGLE. This portion of the analysis develops techniques specific to buyers with restricted complements. In Section 5, we show that BREV gets a 12-approximation to NON-FAVORITE. This portion of the analysis will look somewhat standard to the reader familiar with [6], with a little extra work to extend their main ideas to our setting. We conclude this section with our main theorem, whose proof will be completed by the end of Section 5:

Theorem 3.2. *For a distribution D that has complementary d, $REV \leq (4d{+}16)\max\{BREV, SREV^*\}$.*

[7] We will not actually make use of bullet 3 other than to simplify notation, but it might help remind some readers where these prices come from.

4 Bounding SINGLE

In this section, we prove that the better of selling separately and selling together gets an $O(d)$ approximation to SINGLE.

Proposition 4.1. *SINGLE* $\leq 4d\,SREV^* + 4\,BREV$.

We begin by relating SINGLE to $\text{OPT}^{\text{copies}}$:

Observation 4.2. *SINGLE* $\leq OPT^{copies}$.

Proof. First, observe that there is exactly one S for which $\mathbb{1}[v \in R_S] = 1$. So it is certainly the case that for all v (with $v(S) = \sum_{T \subseteq S} w(T)$), we have:

$$\sum_{S \subseteq M} \max\{0, \bar{\varphi}_S(w(S))\} \cdot \mathbb{1}[v \in R_S] \leq \max_{S \subseteq M}\{0, \bar{\varphi}_S(w(S))\}.$$

$$\Rightarrow \mathbb{E}_{v \leftarrow D}\left[\sum_{S \subseteq M} \max\{0, \bar{\varphi}_S(w(S))\} \cdot \mathbb{1}[v \in R_S]\right] \leq \mathbb{E}_{v \leftarrow D}\left[\max_{S \subseteq M}\{0, \bar{\varphi}_S(w(S))\}\right].$$

Above, the LHS is exactly SINGLE, and the RHS is exactly $\text{OPT}^{\text{copies}}$. □

Note that if the buyer's valuation were additive, at this point we'd already be finished. We could simply set the prices guaranteed by Theorem 2.1 and be done. As we consider more complex buyer valuations, there are two barriers we must overcome. The first is due to substitutability: if we try to set prices on each subset separately, just because the buyer is *willing* to purchase set S doesn't mean he will *choose* to purchase set S, because he may purchase some substitutes instead. Note that this issue doesn't arise in absense of substitutes: if the buyer is willing to purchase S by itself, he is certainly willing to add S to any other set of purchased items. The second barrier is due to complementarity: even once we decide the "correct" price to charge for set S, we can only set prices on *items* and not on *bundles*. Therefore, the prices we want to set for different bundles necessarily interfere with each other. This is the novel barrier unique to values with complementarity, and is also the only part of the analysis where the (necessary) factor of d arises.

The first step to overcoming the complements barrier is to find a subset of bundles for which we can still set the appropriate prices. As a warm-up, let's see what the argument would look like assuming that there were only complements and no substitutes ($\mathcal{C} = 2^M$):

Lemma 4.3. *Let* $\mathcal{C} = 2^M$ *and* $T_1, \ldots, T_k$ *be subsets of* M *such that* $T_i \not\subseteq \cup_{j \neq i} T_j$ *for all* i. *Then for all* $\{p_T\}_{T \subseteq M}$, $SREV \geq \sum_i p_{T_i} \Pr_{x \leftarrow D'_{T_i}}[x \geq p_{T_i}]$.

Proof. Set price p_{T_i} on the item contained in T_i but not $\cup_{j \neq i} T_j$ (if there are multiple, select one arbitrarily). Then by hypothesis, the price the bidder would have to pay in order to receive the entire set T_i is exactly p_{T_i}. Because $\mathcal{C} = 2^M$, whenever $w(T_i) \geq p_{T_i}$, the buyer will choose to purchase the set T_i in addition to whatever else they choose to purchase. Therefore, the item contained in T_i but not $\cup_{j \neq i} T_j$ is purchased with probability at least $\Pr_{x \leftarrow D_{T_i}}[x \geq p_{T_i}]$, and the revenue of this item pricing is at least $\sum_i p_{T_i} \Pr_{x \leftarrow D'_{T_i}}[x \geq p_{T_i}]$. □

The proof of Lemma 4.3 makes use of the assumption that $\mathcal{C} = 2^M$ in exactly one place: to argue that whenever $w(T_i) \geq p_{T_i}$, the buyer chooses to purchase the complete set T_i. When $\mathcal{C} \neq 2^M$, it may be the case that even though the buyer is willing to purchase set T_i, she chooses to purchase substitutes instead. We can remove this assumption on $\mathcal{C}$ by restricting attention to certain price vectors.

Lemma 4.4. *Let $\mathcal{C}$ be any downwards closed set system and $T_1, \ldots, T_k$ be subsets of M such that $T_i \not\subseteq \cup_{j \neq i} T_j$ for all i. Then for all $\{p_T\}_{T \subseteq M}$ such that $p_T \geq 4BREV$ for all T, $SREV^* \geq \frac{1}{4} \sum_i p_{T_i} \Pr_{x \leftarrow D'_{T_i}}[x \geq p_{T_i}]$.*

Proof. Set price $p_{T_i}/2$ on the item contained in T_i but not $\cup_{j \neq i} T_j$ (if there are multiple, again select one arbitrarily). The price the bidder would have to pay in order to receive the entire set T_i is exactly $p_{T_i}/2$. Suppose $w(T_i) \geq p_{T_i}$. Then, the buyer is not only willing to purchase T_i, but also gets utility at least $p_{T_i}/2$ for doing so. The only reason she would choose not to purchase this set is if there were some other set S with $T_i \not\subseteq S$ and $v(S) \geq p_{T_i}/2 \geq 2\text{BREV}$. As $v(S) \leq v(M) - w(T_i)$ for all such S, in order for such a set to exist, it must be the case that $v(M) - w(T_i) \geq 2\text{BREV}$. Clearly, this occurs with probability at most $\frac{1}{2}$, as otherwise we could set price 2BREV on the grand bundle, sell with probability strictly larger than $\frac{1}{2}$ and make revenue strictly larger than BREV. Moreover, $v(M) - w(T_i) = \sum_{U \neq T_i} w(U)$ is completely independent of $w(T_i)$. Therefore, even conditioned on $w(T_i) \geq p_{T_i}$, the probability that the bidder is interested in some other set S with $T_i \not\subseteq S$ is at most $\frac{1}{2}$, and therefore the buyer indeed chooses to purchase T_i with probability at least $\Pr_{x \leftarrow D'_{T_i}}[x \geq p_{T_i}] \cdot \frac{1}{2}$. ∎

Finally, we can combine Lemma 4.4 with Theorem 2.1 to reduce our search to the problem of partitioning the hyperedges into collections $H_x = \{T_{x1}, \ldots, T_{xk_x}\}$ such that $T_{xi} \not\subseteq \cup_{j \neq i} T_{xj}$ for all i.

Corollary 4.5. *Let $\mathcal{C}$ be any downwards closed set system, and let $\{H_x\}_{x \in [k]}$ be a partition of the hyperedges $\{T : f_T(0) < 1\}$ such that for all x, and all $T \in H_x$, $T \not\subseteq \cup_{T' \in H_x \setminus \{T\}} T'$. Then $4kSREV^* + 4BREV \geq SINGLE$.*

Proof. Take $q = 1$ in Theorem 2.1 and let $\{p_T\}_{T \subseteq M}$ be the guaranteed (randomized) prices. By Theorem 2.1 condition 3, there exist two deterministic prices $p_T^H \geq p_T^L$ and probabilities q_T such that $p_T = p_T^H$ with probability q_T, and $p_T = p_T^L$ with probability $1 - q_T$. Therefore, Theorem 2.1 condition 1 can be rewritten as:

$$\text{OPT}^{\text{copies}} \leq \sum_{T \subseteq M} q_T p_T^H \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^H] + (1 - q_T) p_T^L \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^L]$$

We can further rewrite this by breaking up the two sums into prices that exceed 4BREV, and those that don't, let $\mathcal{B} = 4\text{BREV}$ for simplicity:

$$\begin{aligned}\text{OPT}^{\text{copies}} \leq & \sum_{T \subseteq M, p_T^H \leq \mathcal{B}} q_T p_T^H \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^H] + \sum_{T \subseteq M, p_T^L \leq \mathcal{B}} (1 - q_T) p_T^L \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^L] \\ & + \sum_{T \subseteq M, p_T^H > \mathcal{B}} q_T p_T^H \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^H] + \sum_{T \subseteq M, p_T^L > \mathcal{B}} (1 - q_T) p_T^L \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^L]\end{aligned}$$

By condition 2 of Theorem 2.1, we have

$$\sum_{T \subseteq M} q_T \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^H] + (1 - q_T) \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T^L] \leq 1$$

Therefore, as all prices in the top sum above are at most $\mathcal{B}$, the entire top two terms sum to at most $\mathcal{B} = 4\text{BREV}$.

For the bottom two terms, there is no term for T if $p_T^H \leq \mathcal{B}$. If $p_T^H > \mathcal{B} \geq p_T^L$, define $p_T = p_T^H$. If $p_T^H > p_T^L > \mathcal{B}$, then set p_T to whichever of $\{p_T^H, p_T^L\}$ maximizes $p_T \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T]$. Then

$\sum_{T \subseteq M, p_T^H > \mathcal{B}} p_T \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T]$ is at least as large as the bottom two terms above. Moreover, as all $p_T > \mathcal{B}$, we can apply Lemma 4.4 to conclude that for all $T_1, \ldots, T_k$ such that $T_i \not\subseteq \cup_{j \neq i} T_j$ for all i, $\text{SREV}^* \geq 1/4 \sum_i p_{T_i} \Pr_{x \leftarrow D'_{T_i}}[x \geq p_{T_i}]$.

Finally, as $\{H_x\}_{x \in [k]}$ partitions the hyperedges so that for all x and $T \in H_x$, $T \not\subseteq \cup_{T' \in H_x \setminus \{T\}} T'$, we get:

$$\sum_{T \subseteq M, p_T^H > \mathcal{B}} p_T \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T] = \sum_{x=1}^{k} \sum_{T \in H_x, p_T^H > \mathcal{B}} p_T \cdot \Pr_{x \leftarrow D'_T}[x \geq p_T] \leq 4k \cdot \text{SREV}^*$$

The last inequality is due to Lemma 4.4, and completes the proof. □

So the last remaining task is to find a good partition of hyperedges, such that within each partition, every hyperedge contains at least one item not contained in the other hyperedges in the same partition. We isolate this contribution in Section 4.1 below.

4.1 Partitioning Hyperedges with Restricted Complements

Partition-Edges
Input: List of hyperedges, $E \subseteq 2^M$.
Output: A partition of E into $\{H_x\}_x$ such that for all x and all $T \in H_x$, $T \not\subseteq \cup_{T' \in H_x \setminus \{T\}} T'$.

1. $E_{\text{curr}} \leftarrow E$, $i \leftarrow 0$.
2. While $E_{\text{curr}} \neq \emptyset$:
   (a) $i \leftarrow i + 1$
   (b) $E_i \leftarrow E_{\text{curr}}$.
   (c) For each $T \in E_i$ (in arbitrary order): If $T \subseteq \bigcup_{S \in E_i \setminus \{T\}} S$ Then $E_i \leftarrow E_i \setminus \{T\}$.
   (d) $E_{\text{curr}} \leftarrow E_{\text{curr}} \setminus E_i$.
3. Return the partition $\{E_j\}_{j \in [i]}$.

Figure 1: An edge partitioning process.

We provide a high-level description of our algorithm here, and give pseudocode in Figure 1. Recall that the algorithm takes as input a set of hyperedges, and returns a partition of the hyperedges $\{H_x\}_x$, so that in each partition H_x, every hyperedge $S \in H_x$ contains an item that is not in any other hyperedge $T \in H_x$. The algorithm iteratively constructs each H_x, and initially initializes H_x to contain all remaining hyperedges. Then, it iteratively eliminates all "bad" hyperedges (those that *don't* contain an item absent from the others) until the remaining hyperedges have the desired property. In the proof of Theorem 4.6 below, it is easy to show that the algorithm outputs a feasible partition, and the trick is guaranteeing that each iteration makes sufficient progress towards finalizing the partition.

Theorem 4.6. *For any set of hyperedges $E \subseteq 2^M$, Algorithm 1 returns a partition of $E = \{H_x\}_{x \in [k]}$ such that:*

1. *For all x, and all $T \in H_x$, $T \not\subseteq \cup_{T' \in H_x \setminus \{T\}} T'$.*

2. $k \leq \max_i\{|\{T \in E : i \in T\}|\}$.

Proof. First, it is clear that the algorithm indeed properly outputs a partition of E: observe that due to line 2d, when a hyperedge is permanently assigned to some E_i, it will not be assigned to any $E_{i'}$, which implies that all the E_i's are disjoint. Also, every hyperedge is either permanently assigned to some E_i, or remains in E_{curr}, which, by line 2 implies that the algorithm terminates only when every hyperedge is permanently assigned to some E_i. So every hyperedge is contained in some partition, and the partitions are disjoint.

That the output partition satisfies Property 1) is easy to verify: For any x, $T \in H_x$ only the check in 2c passes for T and (the present) H_x. Once the check passes, some other edges will be removed from H_x before the output. Clearly, removing edge from H_x cannot cause T to all of a sudden be contained in $\cup_{T' \in H_x \setminus \{T\}} T'$ when it was previously not contained. So Property 1) is satisfied.

To prove Property 2), first denote by E^i_{curr} the state of E_{curr} at the start of iteration i. We will show that $\cup_{T \in E_i} T = \cup_{T \in E^i_{\text{curr}}} T$. In other words, every element contained in some hyperedge in E^i_{curr} is still contained in some hyperedge in E_i. To see this, observe that when E_i is first set to E^i_{curr}, we clearly have $\cup_{T \in E_i} T = \cup_{T \in E^i_{\text{curr}}} T$. The only time hyperedges are removed from E_i is in step 2c. Note that in order for a hyperedge to be removed from E_i, it must be the case that $T \subseteq \cup_{T' \in E_i \setminus \{T\}} T'$. In other words, in order to remove T from E_i, it must be that all the elements contained in T are also contained in $\cup_{T' \in E_i \setminus \{T\}} T'$. Therefore, removing T does not change $\cup_{T' \in E_i} T'$, and when we terminate, we maintain $\cup_{T \in E_i} T = \cup_{T \in E^i_{\text{curr}}} T$.

To see why this implies Property 2), note that the above implies that if for any i, $|\{T \in E, i \in T\}| = d$, then i will be contained in at least one hyperedge in all of $E_1, \ldots, E_d$, and therefore no hyperedges containing i remain in E^{d+1}_{curr}. In particular, for $d = \max_i\{|\{T \in E, i \in T\}|\}$, it's the case that for all i, no hyperedges containing i remain in E^{d+1}_{curr}, and therefore the algorithm terminates with at most d partitions. □

We can now combine everything to provide a proof of Proposition 4.1:

Proof of Proposition 4.1. Combining Theorem 4.6 with Corollary 4.5, we get that whenever D has complementarity d, that $4d\text{SREV}^* + 4\text{BREV} \geq \text{SINGLE}$, completing the proof. □

5 Bounding NON-FAVORITE

In this section, we bound NON-FAVORITE using similar ideas to those developed in [6]. Much of the process will look familiar to experts famliar with [41, 6], but there are a couple of new ideas sprinkled in. We begin by breaking NON-FAVORITE into CORE + TAIL, as is by now standard (t will be chosen later). Omitted proofs appear in Appendix C.

Lemma 5.1. *NON-FAVORITE is upper bounded by the following:*

$$\mathop{\mathbb{E}}_{v \leftarrow D}\left[\max_{S \in \mathcal{C}}\{\sum_{T \subseteq S} w(T) \cdot \mathbb{1}\left[w(T) \leq t\right]\}\right] + \qquad (\textit{CORE})$$

$$\mathop{\mathbb{E}}_{v \leftarrow D}\left[\sum_{S : w(S) > t} w(S) \cdot \mathbb{1}\left[v \notin R_S\right]\right] \qquad (\textit{TAIL})$$

Bounding CORE. Our main approach to bound CORE is to apply the same concentration bound of Schechtman [42] used in [41]. Essentially, we just have to show that our valuation functions are "subadditive over independent items," for the appropriate definition of "items" (which happens to be hyperedges). It's perhaps not obvious that our valuation functions are subadditive over independent "items," but indeed they are.

Let's first recall the definition of subadditive over independent items. In the definition below, we intentionally write N instead of M to denote the set of items, as the "items" in the definition may be different than the items for sale.

Definition 1. *A distribution D over valuation functions $v : 2^N \to \mathbb{R}$ is* subadditive over independent items *if the following conditions hold:*

1. No externalities and independence across items*: For every item i, let Ω_i be a compact subset of a normed space (i.e., $\Omega_i = [0,1]$). There exists a product distribution D' over $\times_{i\in N}\Omega_i$ (that is, $D' = \prod_{i\in N} D'_i$), and a collection of deterministic functions $V_S : \times_{i\in S}\Omega_i \to \mathbb{R}$ such that a sample v from D can be drawn by sampling $\vec{x} \leftarrow D'$, and defining $v(S) = V_S(\vec{x}_S)$.*
2. Monotonicity*: Every v in the support of D is monotone, i.e., $v(\mathcal{S}) \le v(\mathcal{S}')$ for every $\mathcal{S} \subseteq \mathcal{S}'$.*
3. Subadditivity*: Every v in the support of D is subadditive, i.e., $v(\mathcal{S}\cup\mathcal{S}') \le v(\mathcal{S}') + v(\mathcal{S}')$, for all S, S'.*

Definition 2. *Let D denote a distribution over valuation functions, and D' denote the product distribution and $\{V_S(\cdot)\}$ the deterministic functions that witness D as subadditive over independent items. Then D is c-*Lipschitz *if for all $\vec{x}, \vec{y}$, and sets of items S, T, we have:*

$$|V_S(\vec{x}_S) - V_T(\vec{y}_T)| \le c \cdot (|X \cup Y| - |X \cap Y| + |\{i \in X \cap Y : x_i \neq y_i\}|)$$

We use the following lemma and corollary (of a concentration inequality due to Schechtman [42]) from [41] (the bound in Corollary 5.3 is slightly improved from [41], so we include a proof in Appendix C):

Lemma 5.2. *([41]) Let D be a distribution that is subadditive over independent hyperedges, where for each hyperedge T, $v(\{T\}) \in [0, c]$ with probability 1. Then D is c-Lipschitz.*

Corollary 5.3. *([41]) Suppose that D is a distribution that is subadditive over independent hyperedges and c-Lipschitz, if a is the median of $v(N)$, then $\mathbb{E}[v(N)] \le 3a + c \cdot (2 + 1/\ln 2)$*

Finally, we just need to relate CORE to a random variable that is subadditive over independent items.

Lemma 5.4. *CORE is the expectation of a random variable $v_{CORE}(N)$, where $v_{CORE}(\cdot)$ is t-lipschitz and subadditive over independent items $N = 2^M$. Moreover, $v_{CORE}(N)$ is stochastically dominated by $v(M)$.*

Proof. Let the "items" $N = 2^M$. Let the distributions $\hat{D}_T = D'_T \cdot \mathbb{1}[w(T) \le t]$ (that is, a random variable drawn from $\hat{D}_T$ can be coupled with the random variable $w(T) \cdot \mathbb{1}[w(T) \le t]$). Define constraints $\mathcal{C}' \subseteq 2^N$ $(= 2^{2^M})$ so that a subset U of 2^M is in $\mathcal{C}'$ if and only if there exists a set $C \in \mathcal{C}$ with $\cup_{T\in U} T \subseteq C$. In other words, $U \in \mathcal{C}'$ if and only if the union of elements of U is contained in some set in $\mathcal{C}$. Finally, define $V_U(\vec{x}_U) = \max_{U'\subseteq U, U'\in\mathcal{C}'}\{\sum_{T\in U} x_T\}$.

It is easy to see that $v_{\text{CORE}}(\cdot)$ has no externalities and independent items. It is also easy to see that $v_{\text{CORE}}(\cdot)$ is monotone. Finally, we'll prove that $v_{\text{CORE}}(\cdot)$ is subadditive by observing

that $\mathcal{C}'$ is downwards-closed. To see this, simply observe that if $U' \subseteq U$, and $\cup_{T\in U} T \subseteq C$, then clearly $\cup_{T\in U'} T \subseteq C$. So if $C \in \mathcal{C}$ witnesses that $U \in \mathcal{C}'$ and $U' \subseteq U$, then C also witnesses that $U' \in \mathcal{C}'$.

Now that $\mathcal{C}'$ is downwards closed, it's easy to see (and well-known) that v_{CORE} is subadditive: For any U, W, let $X = \arg\max_{X'\subseteq U\cup W, X\in\mathcal{C}'}\{\sum_{T\in X} x_T\}$. Then let $U' = X \cap U$, and $W' = X \cap W$. Clearly, $\sum_{T\in X} x_T \leq \sum_{T\in U'} x_T + \sum_{T\in W'} x_T$. As $\mathcal{C}'$ is downwards closed, $U' \in \mathcal{C}'$ and $W' \in \mathcal{C}'$. Therefore, $v_{\text{CORE}}(W) + v_{\text{CORE}}(U) \geq \sum_{T\in U'} x_T + \sum_{T\in W'} x_T \geq \sum_{T\in X} x_T = v_{\text{CORE}}(U \cup W)$, and $v_{\text{CORE}}(\cdot)$ is subadditive.

So finally, we just have to show that $v_{\text{CORE}}(N)$ is stochastically dominated by $v(M)$. Couple the random variable x_T drawn from $\hat{D}_T$ so that $x_T = w(T) \cdot \mathbb{1}[w(T) \leq t]$. Now consider $U^* = \arg\max_{U\subseteq 2^M, U\in\mathcal{C}'}\{\sum_{T\in U} x_T\}$. Then we have $v_{\text{CORE}}(N) = \sum_{T\in U^*} x_T$. By definition of $\mathcal{C}'$, there exists some $C \in \mathcal{C}$ such that $T \subseteq C$ for all $T \in U^*$. Therefore:

$$\begin{aligned}
v_{\text{CORE}}(N) = \sum_{T\in U^*} x_T &\leq \sum_{T\subseteq C} x_T \\
&\leq \sum_{T\subseteq C} w(T) && \text{(because } x_T \leq w(T)\text{)} \\
&\leq \max_{S\subseteq M, S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\} && \text{(because } C \in \mathcal{C}\text{)} \\
&= v(M).
\end{aligned}$$

So when x_T and $w(T)$ are coupled in this way, we have $v_{\text{CORE}}(N) \leq v(M)$, and therefore $v(M)$ stochastically dominates $v_{\text{CORE}}(N)$. □

Now, Lemma 5.4 combined with Corollary 5.3 essentially says that $3 \cdot v(M)$ exceeds $\text{CORE} - t \cdot (2 + 1/\ln 2)$ with probability at least $1/2$, allowing us to conclude with the following proposition:

Proposition 5.5. *CORE* ≤ 6*BREV* $+ t \cdot (2 + 1/\ln 2)$.

Proof. Let a be the median of the random variable $v_{\text{CORE}}(N)$. Then $\Pr[v_{\text{CORE}}(N) \geq a] = 1/2$. As $v(M)$ stochastically dominates $v_{\text{CORE}}(N)$, we have $\Pr[v(M) \geq a] \geq 1/2$. Moreover, by Corollary 5.3, the fact that $\text{CORE} = \mathbb{E}[v_{\text{CORE}}(N)]$, and that v_{CORE} is t-lipschitz and subadditive over independent items, we have:

$$\text{CORE} \leq 3a + t(2 + 1/\ln 2).$$

Moreover, as $\Pr[v(M) \geq a] \geq 1/2$, we have:

$$\text{BREV} \geq a/2.$$

Combining the two above equations proves the proposition. □

Bounding TAIL. Our approach to bound TAIL is again similar to [6]. We begin by rewriting TAIL using linearity of expectation and the fact that the hypergraph representation w of valuation

v is drawn from D' which is a product distribution:

$$\begin{aligned}
\text{TAIL} &= \mathop{\mathbb{E}}_{v\leftarrow D}\left[\sum_{T\subseteq M, w(T)>t} w(T)\cdot \mathbb{1}\left[v\notin R_T\right]\right] = \mathop{\mathbb{E}}_{v\leftarrow D}\left[\sum_{T\subseteq M, w(T)>t} w(T)\cdot \mathbb{1}\left[\exists T', w(T')>w(T)\right]\right] \\
&= \sum_{T\subseteq M} \mathop{\mathbb{E}}_{v\leftarrow D}\left[w(T)\cdot \mathbb{1}\left[w(T)>t\wedge v\notin R_T\right]\right] && \text{(by linearity of expectation)} \\
&= \sum_{T\subseteq M}\sum_{x>t, f_T(x)>0} x\cdot f_T(x)\cdot \Pr_{D_{-T}}\left[\exists T', w(T')>x\right] && \text{(by independence across hyperedges)}
\end{aligned}$$

From here, we use essentially the same lemma from [6]. We have replaced their SREV with BREV, but the proof is identical.

Lemma 5.6 ([6])**.** *For all x, T, $x\cdot \Pr_{w\leftarrow D'_{-T}}[\exists T', w(T')>x]\le BREV$.*

Proposition 5.7. *$TAIL\le \left(\sum_{T\subseteq M}\Pr[w(T)>t]\right)\cdot BREV$.*

Setting the Cutoff. Finally, we just need an appropriate choice of t. We'll choose to set t such that $\sum_{T\subseteq M}\Pr[w(T)>t]=k$ for the appropriate choice of k. We first show how to relate t to BREV. Lemma 5.8 below is well-known, but we provide a proof in Appendix C for completeness.

Lemma 5.8. *Let $E_1,\dots,E_k$ be independent events such that $\sum_i \Pr[E_i]=k$. Then $\Pr[\cup_i E_i]\ge 1-e^{-k}$.*

Corollary 5.9. *If t is such that $\sum_{T\subseteq M}\Pr[w(T)>t]=k$, then $BREV\ge(1-e^{-k})t$.*

Proof. Apply Lemma 5.8 to the events $E_T=\{w(T)>t\}$. Then the probability that there exists some hyperedge T with $w(T)>t$ is at least $(1-e^{-k})$. So the grand bundle will sell at price t with probability at least $(1-e^{-k})$. □

We can now complete our bound for NON-FAVORITE, and the proof of Theorem 3.2

Proposition 5.10. *$NON\text{-}FAVORITE\le 12 BREV$*

Proof. Combine Propositions 5.5 and 5.7 taking t such that $\sum_T \Pr[w(T)>t]=1.66$. □

Proof of Theorem 3.2. Simply combine Propositions 4.1 and 5.10 with Corollary 3.1. □

6 Lower bounds

The following proposition shows that the factor d approximation (established in Theorem 3.2) is tight (up to constant factors), even when there are no substitutes ($\mathcal{C}=2^M$).

Proposition 6.1. *There exists a distribution D with complementarity d, for which $REV\ge\frac{d}{4}\max\{BREV, SREV\}$.*

Furthermore, we argue that this parameter correctly characterizes the degree of complementarity in our setting. Specifically, in Proposition 6.2, we establish extremely high lower bounds (as a function of the complementarity degree) on the approximation ratio that can be obtained by $\max\{\text{BREV}, \text{SREV}\}$ for previous measures of complementarity from the literature. In what follows

we give informal definitions of the different measures of complementarities and state their lower bounds. Formal definitions are deferred to Appendix D.

A valuation is in PH-k [1] if its hypergraph representation w has only positive hyperedges S of size at most k. The supermodular degree (SM) [21], roughly, measures the distance of a valuation from being submodular; it ranges between 1 to m. A valuation is in PS-k if in its hypergraph representation every item shares a positive hyperedge with at most k other items. It can be shown that PS-$k \subseteq$ SM-k, thus every lower bound on PS-k carries over to SM-k. Also, these lower bounds trivially hold for "maximum over PH" [20] and "maximum over PS" [22] hierarchies. The following proposition asserts the lower bounds for the aforementioned hierarchies.

Proposition 6.2. *The following hold for distributions in our settings, where hyperedges values $w(T)$ are independently drawn, and $v(S) = \sum_{T \subseteq S} w(T)$.*

1. *There exists a distribution D with only PH-k valuations in the support, for which $REV \geq \frac{1}{2m} \sum_{1 \leq i \leq k} \binom{m}{i} \max\{BREV, SREV\}$. E.g., for PH-2, $REV \geq \Omega(m) \cdot \max\{BREV, SREV\}$.*
2. *There exists a distribution D with only PS-k valuations in the support, for which $REV \geq \frac{2^{k+1}-1}{2(k+1)} \max\{BREV, SREV\}$.*

Consider a set of hyperedges E (to be defined per-case). Index the hyperedges with integers in $\{1+a, 2+a, \ldots, |E|+a\}$ (we abuse notation and use e both for index and hyperedge, i.e., set of items). The product distribution D' has $f_e(0) = 1$ for all $e \notin E$, and for every $e \in E$, set $f_e(0) = 1 - 2^{-e}$, and $f_e(2^e) = 2^{-e}$. Let D be the distribution that samples $w \leftarrow D'$ and returns $v(S) = \sum_{T \subseteq S} w(T)$.

Proposition 6.3. *For the above distribution D, we have $REV(D) \geq |E|$, but $SREV(D) \leq 2m$ and $BREV(D) \leq 2$.*

Proof. First, consider the random variable $v(M)$. We have $v(M) \leq \sum_{e=1+a}^{|E|+a} w(e)$. For any price p, in order to have $v(M) \geq p$, we must have $w(e) > 0$ for some $e \geq \log p$, as $\sum_{e=1+a}^{\log p - 1} 2^e = p - 2^{a+1} < p$.[8] Note that the there is no reason to price below 2^{1+a}. But also, by union bound, the probability that this occurs is at most $\sum_{e \geq \log p} 2^{-e} \leq 2^{1-\log p} \leq 2/p$. So for any price p we could set on the grand bundle, it sells with probability at most $2/p$, so BREV ≤ 2.

Similarly, for any price p_i, in order for the buyer to possibly be willing to purchase item i, we must have $\sum_{e \ni i} w(e) \geq p_i$. Again, in order for this to happen, we must have $w(e) > 0$ for some $e \geq \log p_i$, $e \ni i$. And again by union bound, the probability that this occurs is at most $2/p_i$. So for any price p_i we could set on item i, the probability that the buyer is possibly willing to purchase item i is at most $2/p_i$, so SREV $\leq 2m$.

Consider however the following mechanism, which essentially sells the hyperedges in E separately. The mechanism allows the buyer to purchase any set S she chooses, and charges price 2^S. By union bound,[9] the probability that $v \equiv 0$ is at least $1 - 2^{-a}$. Therefore, whenever $w(e) > 0$, with probability at least $1 - 2^{-a}$, the buyer will choose to purchase exactly the set e and pay 2^e. So the revenue is at least $\sum_{e=1+a}^{|E|+a} 2^{-e} \cdot 2^e \cdot (1 - 2^{-a}) = |E| \cdot (1 - 2^{-a})$. Taking $a \to \infty$ completes the proof. ☐

Let us now see how proposition 6.3 implies proposition 6.1 and proposition 6.2.

[8] $\sum_{e=1+a}^{n-1} 2^e = 2^n - 2^{a+1}$

[9] $\sum_{e=1+a}^{n+a} 2^{-e} = 2^{-a} - 2^{-n-a}$ therefore its complement is at least $1 - 2^a$

proof of proposition 6.1. Consider a d regular graph (M, E) over m nodes. By definition, every node is contained in exactly d hyperedges. Therefore, if E is the set of hyperedges used to construct D prior to proposition 6.3, then D has complementarity d, and $|E| = md/2$. □

proof of proposition 6.2. To show 6.2.1, consider the set of all hyperedges of size at most k, and apply proposition 6.3. To show proposition 6.2.2, assume for simplicity that m is divisible by $k+1$. Partition M to $m/(k+1)$ sets $M_1, M_2, \dots M_{m/(k+1)}$, all of size $k+1$, and consider all hyperedges $S \subseteq M_i$ for all i. Every item i in M_j has neighbors only from M_j, therefore every valuation in the support is from PS-k. The number of hyperedges is $\frac{m}{k+1} \cdot (2^{k+1} - 1)$. Applying proposition 6.3 completes the proof. □

A Background on Duality Framework

We first recall the duality approach of [6]:

Definition 3. *[Reworded from [6], Definitions 2 and 3] A mapping $\lambda : V \times V \to \mathbb{R}^+$ is* flow-conserving *if for all $v \in V$: $\sum_{v' \in V} \lambda(v, v') \leq f(v) + \sum_{v' \in V} \lambda(v', v)$.[10] The* virtual transformation *associated with λ, Φ^λ, is a transformation from valuation functions in V to valuation functions in $V^\times$ (the closure of V under linear combinations) and satisfies:*[11]

$$\Phi^\lambda(v)(\cdot) = v(\cdot) - \frac{1}{f(v)} \sum_{v' \in V} \lambda(v', v)(v'(\cdot) - v(\cdot)).$$

In the above definition, one should interpret $\lambda(\cdot, \cdot)$ as being potential Lagrangian multipliers for incentive constraints in a certain LP to find the revenue-optimal mechanism, and think of $f(v)$ flow going into each v from some super source, $\lambda(v, v')$ flow going from v to v', and all excess flow (that enters v but doesn't leave) as going from v to a super sink. Note that whether or not a given λ is flow-conserving depends on the population D. Cai et al. show that Lagrangian multipliers that satisfy the above flow conservation constraint yield upper bounds of the following form.

Theorem A.1. *[Reworded from [6], Theorem 10] Let $\mathcal{M}$ be any truthful mechanism where a bidder with type v receives items $X(v)$ and pays $p(v)$. Then for all flow-conserving λ, the expected revenue of $\mathcal{M}$ is upper bounded by its expected virtual welfare with respect to λ. That is:*

$$\mathop{\mathbb{E}}_{v \leftarrow D}[p(v)] \leq \mathop{\mathbb{E}}_{v \leftarrow D}\left[\Phi^\lambda(v)(X(v))\right].$$

As an immediate corollary, we can obtain the following upper bound on the revenue of any truthful mechanism by observing that the bound in Theorem A.1 is maximized when $X(v)$ is deterministically $\arg\max_{S \subseteq M}\{\Phi^\lambda(v)(S)\}$.

Corollary A.2. *For all D, and all flow-conserving λ, we have:*

$$REV(D) \leq \mathop{\mathbb{E}}_{v \leftarrow D}\left[\max_{S \subseteq 2^M} \Phi^\lambda(v)(S)\right].$$

[10] This is equivalent to stating that there exists a $\lambda(v, \perp) \geq 0$ such that $\lambda(v, \perp) + \sum_{v' \in V} \lambda(v, v') = f(v) + \sum_{v' \in V} \lambda(v', v)$, which might look more similar to the wording of Definition 2 in [6].

[11] That is, $\Phi^\lambda(v)$ is a (possibly negative) function from 2^M to $\mathbb{R}$, and satisfies $\Phi^\lambda(v)(S) = v(S) - \frac{1}{f(v)} \sum_{v' \in V} \lambda(v', v)(v'(S) - v(S))$. for all $S \subseteq M$.

We begin this section by defining our flow-conserving λ and the resulting Φ^λ. Readers familiar with [6] will recognize it as the natural generalization of their flow to our setting, and we will make the language as similar as possible.

We will break V into 2^m-1 different regions, depending on which hyperedge is the most valuable to a buyer with value v. Specifically, we say that v is in region R_A if $A = \arg\max_{T\subseteq M}\{w(T)\}$, with ties broken lexicographically. Recall that D is established by drawing w from the product distribution D' and the returned valuation v satisfies $v(S) = \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T} w(U)\}$. Then consider the following flow:

Definition 4 (Flow for our benchmark)**.** *If $v \in R_A$, define $w'(T) = w(T)$, for all $T \neq A$, and define $w'(A) = \min_{x>w(A)}\{x : f_A(x) > 0\}$. Set $\lambda(v', v) = \Pr_{x\leftarrow D'_A}[x \geq w'(A)] \cdot \prod_{T\neq A} f_T(w'(T)) = f(v) \cdot \frac{\Pr_{x\leftarrow D'_A}[x\geq w'(A)]}{f_A(w(A))}$ for the $v'(\cdot)$ such that $v'(S) = \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T} w'(U)\}$ for all S, and $\lambda(v'', v) = 0$ for all other v''.*

Proposition A.3. *The $\lambda(\cdot,\cdot)$ from Definition 4 is flow-conserving. Moreover, if $v(\cdot)$ is such that $v(S) = \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T} w(U)\}$, and $v \in R_A$, then Φ^λ satisfies the following:*

$$\Phi^\lambda(v)(S) \leq \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} + \max\{0, \varphi_A(w(A))\} \leq \max_{T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} + \max\{0, \varphi_A(w(A))\}.$$

Proof. That $\lambda(\cdot,\cdot)$ is flow-conserving is clear: every $v \in R_A$ has total incoming flow of $f(v) \cdot \frac{\Pr_{x\leftarrow D'_A}[x\geq w(A)]}{f_A(w(A))}$ ($f(v)$ of this comes from the source, the remaining $f(v)\cdot\frac{\Pr_{x\leftarrow D'_A}[x>w(A)]}{f_A(w(A))}$ comes from other types in R_A). Every $v \in R_A$ also has outgoing flow either equal to 0 (if decreasing the value of $w(A)$ moves the resulting v' out of R_A), or exactly $f(v)\cdot\frac{\Pr_{x\leftarrow D'_A}[x\geq w(A)]}{f_A(w(A))}$ (otherwise). In either case, the flow out is at most the flow in.

Let's now compute $\Phi^\lambda(v)(S)$. Plugging into Definition 3, we get:

$$\Phi^\lambda(v)(S) = v(S) - \frac{(v'(S) - v(S))\Pr_{x\leftarrow D'_A}[x \geq w(A)]}{f_A(w(A))}.$$

Recall that $v'(S) \geq v(S)$ for all S, and therefore $\Phi^\lambda(v)(S) \leq v(S)$ for all S. Now there are two cases to consider: In the first case, maybe $\max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} = v(S)$. In other words, the set in $\mathcal{C}$ "chosen" by a consumer with valuation v doesn't contain A. In this case, we immediately get that $\Phi^\lambda(v)(S) \leq v(S) = \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\}$, as desired.

In the second case, maybe $\max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} < v(S)$. In other words, the set in $\mathcal{C}$ "chosen" by a consumer with valuation v contains A. In this case, increasing $w(A)$ by any $x > 0$ increases $v(S)$ by exactly x. Therefore, we have $v'(S) = v(S) + w'(A) - w(A)$, and therefore:

$$\begin{aligned}
\Phi^\lambda(v)(S) &= v(S) - \frac{(w'(A) - w(A))\Pr_{x\leftarrow D'_A}[x \geq w(A)]}{f_A(w(A))} \\
&= \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T} w(U)\} - \frac{(w'(A) - w(A))\Pr_{x\leftarrow D'_A}[x \geq w(A)]}{f_A(w(A))} \\
&\leq \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} + w(A) - \frac{(w'(A) - w(A))\Pr_{x\leftarrow D'_A}[x \geq w(A)]}{f_A(w(A))} \\
&= \max_{T\subseteq S, T\in\mathcal{C}}\{\sum_{U\subseteq T, U\neq A} w(U)\} + \varphi_A(w(A)).
\end{aligned}$$

The last line uses the definition $\varphi_A(w(A)) = w(A) - \frac{(w'(A)-w(A))\Pr_{x\leftarrow D'_A}[x\geq w(A)]}{f_A(w(A))}$, which may seem unfamiliar to readers more familiar with virtual values for continuous distributions. Indeed, this is the right generalization of Myerson's $\varphi(\cdot)$ for continuous distributions to the discrete setting, and we refer the interested reader to Section 4 of [6] for more discussion. □

Ironing. The astute reader will notice that when D'_S is irregular, the bound we probably want above would replace $\varphi_A(\cdot)$ with $\bar{\varphi}_A(\cdot)$. [6] shows how to design a flow that accomplishes this essentially by adding cycles to λ between adjacent types to "iron out" any non-monotonicities, but for their setting of additive buyers. The exact same approach will work here. We omit a proof and refer the reader to [6] for more detail. This allows us to prove corollary 3.1.

proof of corollary 3.1. Simply combine Corollary A.2 and Proposition A.3, after replacing $\varphi(\cdot)$ in Proposition A.3 with $\bar{\varphi}(\cdot)$. □

B Missing Proofs From Section 2 and 3

Definition 5. *A Random variable X is first-order stochastically dominated (FOSD) by random variable Y if for every x, $\Pr[X \geq x] \leq \Pr[Y \geq x]$.*
Note: *If X is FOSD by Y then $\mathbb{E}[X] \leq \mathbb{E}[Y]$.*

proof of theorem 2.1. Let $\mathbb{1}_q$ be an independent indicator random variable that equals 1 with probability q. Let $\{X_S\}_S$ be non-negative independent random variables that are drawn from the independent distributions.

Consider a tie breaking rule among the sets, and let the event $X_S = \max_T\{X_T\}$ be true only when S also wins in the tie breaking rule. Set $q_S = \Pr[X_S = \max_T\{X_T\}]$. So $\sum_S q_S = 1$. set t_S s.t. $\Pr[X_S \geq t_S] = q \cdot q_S$.

Let us see that the random variable $\mathbb{1}_q \cdot X_S \cdot \mathbb{1}[X_S = \max\{X_T\}]$ is FOSD (definition 5) by $X_S \cdot \mathbb{1}[X_S \geq t_S]$.

For every $x \geq t_S$, it holds that $\Pr[\mathbb{1}_q \cdot X_S \cdot \mathbb{1}[X_S = \max\{X_T\}] \geq x] \leq q\Pr[X_S \geq x]$, while $\Pr[X_S \cdot \mathbb{1}[X_S \geq t_S] \geq x] = \Pr[X_S \geq x]$.

For every $x < t_S$, it holds that $\Pr[\mathbb{1}_q \cdot X_S \cdot \mathbb{1}[X_S = \max\{X_T\}] \geq x] \leq q \cdot q_S$ by definition of q_S, while $\Pr[X_S \cdot \mathbb{1}[X_S \geq t_S] \geq x] = \Pr[X_S \geq t_S] = q \cdot q_S$ by definition of t_S. We get that:

$$\begin{aligned}
\mathbb{E}\left[\max_S\{X_S\}\right] &= \sum_S \mathbb{E}[X_S \cdot \mathbb{1}[X_S = \max\{X_T\}]] \\
&= \frac{1}{q}\sum_S \mathbb{E}[\mathbb{1}_q \cdot X_S \cdot \mathbb{1}[X_S = \max\{X_T\}]] \\
&\leq \frac{1}{q}\sum_S \mathbb{E}[X_S \cdot \mathbb{1}[X_S \geq t_S]]
\end{aligned}$$

Let X_S be the random variable that first draws $x \leftarrow D'_S$ and returns $\max\{0, \varphi_S(x)\}$. Assume the distributions are regular, and refer to [9] for the irregular case. As $t_S \geq 0$ we get:

$$\mathop{\mathbb{E}}_{x\leftarrow D'_S}\left[\max_S\{\varphi_S(x), 0\}\right] \leq \sum_S \mathop{\mathbb{E}}_{x\leftarrow D'_S}[\varphi_S(x) \cdot \mathbb{1}[\varphi_S(x) \geq t_S]]$$

Observe that the above term for each S is the expected virtual value of the mechanism that allocates to a bidder with value x if x exceeds $p_S = \inf\{x : \varphi_S(x) = t_S\}$. This allocation is achieved by posting a price p_S. By Myerson's payment identity:

$$\mathop{\mathbb{E}}_{x\leftarrow D'_S}\left[\varphi_S(x)\cdot\mathbb{1}\left[\varphi_S(x)\geq t'_S\right]\right] = \mathop{\mathbb{E}}_{x\leftarrow D'_S}\left[p_S\cdot\Pr\left[x\geq p_S\right]\right]$$

This concludes property 1. Property 2 follows by monotonicity of φ_S (regularity of D'_S, for the irregular case refer to [9]):

$$\sum_S \Pr_{D'_S}\left[x\geq p_S\right] = \sum_S \Pr_{D'_S}\left[\varphi_S(x)\geq t_S\right] = \sum_S q\cdot q_S = q$$

□

Proof of Theorem 3.2. Simply combine Propositions 4.1 and 5.10 with Corollary 3.1, to get:

$$\mathrm{REV}(D)\leq 4d\mathrm{SREV}^*(D)+16\mathrm{BREV}(D)\leq(4d+16)\max\{\mathrm{SREV}^*,\mathrm{BREV}\}$$

□

C Missing Proofs From Section 5

proof of lemma 5.1. The proof follows from the following algebra:

$$\begin{aligned}
(\text{NON-FAVORITE}) &= \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[v\notin R_T\right]\}\right]\\
&= \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[w(T)\leq t\right]\cdot\mathbb{1}\left[v\notin R_T\right]+w(T)\cdot\mathbb{1}\left[w(T)>t\right]\cdot\mathbb{1}\left[v\notin R_T\right]\}\right]\\
&\leq \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[w(T)\leq t\right]\}\right]\\
&\qquad + \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[w(T)>t\right]\cdot\mathbb{1}\left[v\notin R_T\right]\}\right]\\
&\leq \mathop{\mathbb{E}}_{v\leftarrow D}\left[\max_{S\in\mathcal{C}}\{\sum_{T\subseteq S} w(T)\cdot\mathbb{1}\left[w(T)\leq t\right]\}\right]\\
&\qquad + \mathop{\mathbb{E}}_{v\leftarrow D}\left[\sum_{T|w(T)>t} w(T)\cdot\mathbb{1}\left[v\notin R_T\right]\}\right]
\end{aligned}$$

□

proof of corollary 5.3. By corollary 12 in [42], we know that for all $k>0$:

$$\Pr\left[v(N)\geq 3\cdot a+k\cdot c\right]\leq\min\{1,4\cdot 2^{-k}\} \tag{1}$$

Substituting $x = 3 \cdot a + k \cdot c$ gets $k = (x - 3a)/c$. Therefore, equation (1) becomes meaningful only when $4 \cdot 2^{-k} \leq 1$, i.e., when $x \geq 2c + 3a$. Computing the expected value of $v(N)$ gives:

$$\int_0^\infty \Pr\left[v(N) > x\right] dx \leq \int_0^\infty \min\{1, 4 \cdot 2^{(3a-x)/c}\} dx = 2c + 3a + 4 \cdot 2^{3a/c} \cdot \int_{2 \cdot c + 3a}^\infty 2^{-x/c} \cdot dx$$

Computing the integral gives: $-\frac{c}{\ln 2}\left[2^{-x/c}\right]_{2c+3a}^{\infty} = \frac{c}{\ln 2} \cdot 2^{-\frac{2c+3a}{c}} = \frac{c}{4 \ln 2} \cdot 2^{-3a/c}$, which, plugged back to the equation concludes that

$$\mathbb{E}\left[v(N)\right] \leq 2c + 3a + \frac{c}{\ln 2}$$

as desired. □

proof of lemma 5.6. For any x, we can set price x on the grand bundle. It will sell with probability at least $\Pr_{w \leftarrow D'_{-T}}[\exists T', w(T') > x]$, as whenever there is a single hyperedge with contribution x, certainly the buyer's value for the grand bundle is at least x. Therefore, BREV $\geq x \cdot \Pr_{w \leftarrow D'_{-T}}[\exists T', w(T') > x]$. □

proof of proposition 5.7. By Lemma 5.6, We get:

$$\sum_{T \subseteq M} \sum_{x > t, f_T(x) > 0} x \cdot f_T(x) \cdot \Pr_{w \leftarrow D'_{-T}}\left[\exists T', w(T') > x\right] \leq \sum_{T \subseteq M} \sum_{x > t, f_T(x) > 0} f_T(x) \cdot \text{BREV} = \sum_{T \subseteq M} \Pr[w(T) > t] \cdot \text{BREV}.$$

□

proof of lemma 5.8. By independence:

$$\Pr\left[\cup_i E_i\right] = 1 - \prod_i \left(1 - \Pr\left[E_i\right]\right)$$

So if we define $q_i = \Pr[E_i]$, we want to maximize $\prod_i (1 - q_i)$ subject to $\sum_i q_i = k$. Using a Lagrangian multiplier of λ on the constraint $\sum_i q_i = k$, we get a new objective of:

$$\prod_i (1 - q_i) + \lambda \cdot (\sum_i q_i) - \lambda k$$

We see that the partial with respect to q_i of the above is exactly $-\prod_{j \neq i}(1 - q_j) + \lambda$. So setting $q_i = k/n$ for all i, and $\lambda = (1 - k/n)^{n-1}$, we see that $\sum_i q_i = k$ and the partial of the Lagrangian with respect to q_i is 0 for all i. Therefore, this is the optimal solution. At $q_i = k/n$ for all i, we have $\prod_i (1 - k/n) = (1 - k/n)^n \leq e^{-k}$. □

D Lower Bounds

Below We formally define the previously mentioned complementarity measures.

Definition 6. *[1, 20](MPH) A valuation v is positive hypergraph (PH) of degree at most k if there exists a hyperedge weight function $w \geq 0$, $w(T) = 0$ for all $|T| > k$, so that $v(S) = \sum_{T \subseteq S} w(T)$.*

A valuation v is maximum over PH (MPH) of degree at most k if there exists a collection L of such hyperedge weight functions, so that $v(S) = \max_{\ell \in L}\{\sum_{T \subseteq S} w_\ell(T)\}$.

Definition 7. *[22](MPS) A valuation v is positive supermodular (PS) of degree at most k if there exists a hyperedge weight function $w \geq 0$, so that for every item i, it holds that $|\{j \in M : \exists T, w(T) > 0, \{j, i\} \subseteq T\}| \leq k$, i.e., item i has at most k neighbors (other items that share a positive hyperedge with it).*

A valuation v is maximum over PS (MPS) of degree at most k if there exists a collection L of such hyperedge weight functions, so that $v(S) = \max_{\ell \in L}\{\sum_{T \subseteq S} w_\ell(T)\}$.

Definition 8. *[21](SM) A valuation v is supermodular (SM) of degree at most k if for each item i, the number of items i' so that there exists a set $S_{i'} \not\ni i$ so that $v(S_{i'} \cup i) - v(S_{i'}) > v(S_{i'} \setminus \{i'\} \cup \{i\}) - v(S_{i'} \setminus \{i'\})$ is at most k, i.e., i's marginal contribution to a set may* increase *by adding another item, to at most k different items.*

References

[1] I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 3–16. ACM, 2012.

[2] M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 21–30. IEEE, 2014.

[3] M. Bateni, S. Dehghani, M. Hajiaghayi, and S. Seddighin. Revenue maximization for selling multiple correlated items. In *Algorithms-ESA 2015*, pages 95–105. Springer, 2015.

[4] X. Bei and Z. Huang. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 720–733. SIAM, 2011.

[5] P. Briest, S. Chawla, R. Kleinberg, and S. M. Weinberg. Pricing randomized allocations. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 585–597. Society for Industrial and Applied Mathematics, 2010.

[6] Y. Cai, N. R. Devanur, and S. M. Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 926–939. ACM, 2016.

[7] Y. Cai and M. Zhao. Simple Mechanisms for Subadditive Buyers via Duality. *Manuscript*, 2016.

[8] S. Chawla, J. D. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 243–251. ACM, 2007.

[9] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.

[10] S. Chawla, D. L. Malec, and B. Sivan. The power of randomness in bayesian optimal mechanism design. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 149–158. ACM, 2010.

[11] S. Chawla and J. B. Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 579–596, New York, NY, USA, 2016. ACM.

[12] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 269–286. ACM, 2013.

[13] C. Daskalakis, A. Deckelbaum, and C. Tzamos. The complexity of optimal mechanism design. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1302–1318. Society for Industrial and Applied Mathematics, 2014.

[14] C. Daskalakis and S. M. Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 370–387. ACM, 2012.

[15] N. Devanur, J. Morgenstern, V. Syrgkanis, and S. M. Weinberg. Simple auctions with simple strategies. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 305–322. ACM, 2015.

[16] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103. Springer, 2007.

[17] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 610–618. ACM, 2005.

[18] A. Eden, M. Feldman, O. Friedler, I. Talgam-Cohen, and S. M. Weinberg. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. *Manuscript*, 2016.

[19] U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.

[20] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[21] U. Feige and R. Izsak. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 247–256. ACM, 2013.

[22] M. Feldman, O. Friedler, J. Morgenstern, and G. Reiner. Simple mechanisms for agents with complements. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 251–267, 2016.

[23] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 201–210. ACM, 2013.

[24] M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135. SIAM, 2015.

[25] M. Feldman and R. Izsak. Constrained monotone function maximization and the supermodular degree. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 28. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.

[26] M. Feldman and R. Izsak. Building a good team: Secretary problems and the supermodular degree. *arXiv preprint arXiv:1507.06199*, 2015.

[27] S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 656–656. ACM, 2012.

[28] S. Hart and N. Nisan. The menu-size complexity of auctions. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 565–566. ACM, 2013.

[29] S. Hart and P. Reny. *Maximal revenue with multiple goods: Nonmonotonicity and other observations*. 2012.

[30] J. D. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility via matchings. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 734–747. SIAM, 2011.

[31] J. D. Hartline and B. Lucier. Bayesian algorithmic mechanism design. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 301–310. ACM, 2010.

[32] R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136. ACM, 2012.

[33] D. Lehmann, L. I. Oćallaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49(5):577–602, 2002.

[34] X. Li and A. C.-C. Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[35] J. Morgenstern. *Market Algorithms: Incentives, Learning and Privacy*. PhD thesis, Stanford University, 2015.

[36] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[37] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.

[38] G. Pavlov. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics*, 11(3), 2011.

[39] J. Riley and R. Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, pages 267–289, 1983.

[40] J.-C. Rochet and P. Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, July 1998.

[41] A. Rubinstein and S. M. Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.

[42] G. Schechtman. Concentration, results and applications. *Handbook of the geometry of Banach spaces*, 2:1603–1634, 2003.

[43] J. Thanassoulis. Haggling over substitutes. *Journal of Economic theory*, 117(2):217–245, 2004.

[44] A. C.-C. Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. SIAM, 2015.